



XP 000032477

8277 ACM Transactions on Computer Systems
6 (1988) Febr., No. 1, New York, NY, USA

G06F15/408


G06F3/16

E

p. 3-27

Managing Stored Voice in the Etherphone System

DOUGLAS B. TERRY and DANIEL C. SWINEHART
Xerox Palo Alto Research Center



The *voice manager* in the Etherphone system provides facilities for recording, editing, and playing stored voice in a distributed personal-computing environment. It provides the basis for applications such as voice mail, annotation of multimedia documents, and voice editing using standard text-editing techniques. To facilitate sharing, the voice manager stores voice on a special voice file server that is accessible via the local internet. Operations for editing a passage of recorded voice simply build persistent data structures to represent the edited voice. These data structures, implementing an abstraction called *voice ropes*, are stored in a server database and consist of lists of intervals within voice files. Clients refer to voice ropes solely by reference. *Interests*, additional persistent data structures maintained by the server, serve two purposes: First, they provide a sort of directory service for managing the voice ropes that have been created. More importantly, they provide a reliable reference-counting mechanism, permitting the garbage collection of voice ropes that are no longer needed. These interests are grouped into classes; for some important classes, obsolete interests can be detected and deleted by a class-specific algorithm that runs periodically.

Categories and Subject Descriptors: C.2.4 [Computer-Communication Networks]: Distributed Systems; D.4.2 [Operating Systems]: Storage Management—*allocation/deallocation strategies; storage hierarchies*; D.4.3 [Operating Systems]: File Systems Management; D.4.6 [Operating Systems]: Security and Protection—*access controls; cryptographic controls*; E.2 [Data]: Data Storage Representations; H.2.8 [Database Management]: Database Applications; H.4.3 [Information Systems Applications]: Communications Applications—*electronic mail*

General Terms: Design, Management, Performance, Security

Additional Key Words and Phrases: Digitized voice, network services, storage reclamation, voice-annotated documents, voice editing, voice file server

1. INTRODUCTION

Voice is an important and widely used medium for interpersonal communication. Computers facilitate interpersonal communication through electronic mail and shared documents. Yet our computer systems have traditionally forced us to communicate textually. A major focus of the Etherphone¹ system developed at the Xerox Palo Alto Research Center was to allow voice to be incorporated into computing environments and used in much the same way as text. This paper

¹ Etherphone is a trademark of the Xerox Corporation.

Authors' address: Computer Science Laboratory, Xerox Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1988 ACM 0734-2071/88/0200-0003 \$01.50

ACM Transactions on Computer Systems, Vol. 6, No. 1, February 1988, Pages 3-27.

Best Available Copy

4 • D. B. Terry and D. C. Swinehart

addresses the problems associated with managing stored voice in a distributed computing environment.

The voice-management facilities of the Etherphone system were designed with the following goals in mind:

- Unrestricted use of voice in client applications.* As with text, we want the ability to incorporate voice easily into electronic-mail messages, voice-annotated documents, user interfaces, and other interactive applications. Nicholson gives a good discussion of many office applications that are made possible by treating voice as data [17].
- Sharing among various clients.* If stored voice is to be used as a means of interpersonal communication, then it must be possible for users to easily share a voice passage with one or more colleagues. Clients should be able to share voice as freely as they share files.
- Editing of voice by programs.* Clients should be able to combine previously recorded voice in various ways and insert fresh voice into existing voice passages. The system should permit programmer control over all of these functions.
- Integration of diverse workstations into the system.* Our environment contains a diversity of workstations and associated operating systems that must access the voice-management facilities. Requiring users to learn and adapt to facilities that cannot be well integrated into their existing workplaces is unacceptable.
- Security at least as good as that of conventional file servers.* People are rightfully concerned about the privacy of their communications; the system should take all means to protect this privacy.
- Automatic reclamation of storage occupied by unneeded voice.* Requiring clients to explicitly delete voice passages when they are no longer needed places an unwarranted burden on users and hinders sharing in a distributed system. The system itself should aid in the automatic reclamation of voice storage.

Many of these features are common in traditional file servers that store text files [21]. The characteristics of voice, however, differ significantly from those of text. Standard telephone-quality uncompressed voice occupies 64 kbits of storage per second of recorded voice. This is several orders of magnitude greater than the equivalent typed text. Voice also requires special devices for recording and playing it; that is, a user cannot simply type in a voice passage. More importantly, voice transmission has stringent real-time requirements. These differences dictate special methods for manipulating and sharing voice.

An abstraction that we call *voice ropes* serves as the basis of application-independent methods for recording, playing, editing, and otherwise manipulating digitized voice. The major technical contributions described in this paper involve the use of simple databases to

- (1) describe the results of editing operations such that existing voice passages need not be moved, copied, or decrypted; and
- (2) provide a modified style of reference counting required to allow the automatic reclamation of obsolete voice.

Managing Stored Voice in the Etherphone System • 5




Fig. 1. A simple Etherphone system environment.

The next section provides the background for the rest of this paper, including a quick overview of the Etherphone system architecture as well as some sample voice applications. Section 3 presents the operations on voice ropes designed to support these and future application programs. Section 4 then discusses the design and implementation of the voice rope operations, and the rationale governing various design choices. Section 5 relates our experiences thus far with incorporating voice into workstation applications and examines the system's performance. Section 6 contrasts related work. We conclude by reviewing our design goals and how they were met.

2. BACKGROUND

2.1 The Etherphone System

The Etherphone system is intended for use in a locally distributed computing environment containing multiple workstations and programming environments, multiple networks and communication protocols, and perhaps even multiple telephone transmission and switching choices. The system is intended to be extensible in that introducing new applications, network services, workstations, networks, and other components is possible.

Figure 1 depicts the basic components of the Etherphone system in a simple configuration [23]. Each personal workstation is associated with, but not directly attached to, a microprocessor-based telephone instrument called an *Etherphone*. Etherphones digitize, packetize, and encrypt telephone-quality voice, and transmit it directly over an Ethernet. A *voice manager*, the main topic of this paper, provides storage for voice, telephone conversations, music, and other sounds, recorded at reasonable fidelity. The system also includes other specialized sources or sinks of voice, such as a *text-to-speech server* that receives text strings and returns the equivalent spoken text to the user's Etherphone.

6 • D. B. Terry and D. C. Swinehart

A *voice control server* provides control functions similar to those of a conventional business telephone system and manages the interactions between all other components. In particular, it allows voice-carrying *conversations* to be established rapidly among two or more Etherphones or voice services. An Etherphone conversation is represented by a conversation identifier, a *ConversationID*. The *ConversationID* is distributed by the control server to all participants in the conversation, including Etherphones, workstations, and voice services, when the conversation is established. It is used to identify the conversation in requests and reports issued by the server and participants.

All of the communication required for control in the voice system, such as conversation establishment and the distribution of encryption keys used in voice transmission, is accomplished via a secure remote procedure call (RPC) protocol [4, 5]. Multiple implementations of the RPC mechanisms permit the integration of workstation programs and voice applications programmed in different environments. During the course of a conversation, reports emanating from the voice control server via RPC inform participants about various activities concerning the conversation.

Active parties in a conversation exchange voice using a specialized *voice transmission protocol* [22]. During each conversation all transmitted voice is encoded using DES electronic-codebook (ECB) encryption [16] with a randomly generated encryption key issued by the voice control server.

Workstations are the key to providing enhanced user interfaces and control over the voice capabilities. The extensibility of the local programming environment—be it Cedar, Interlisp, or the Xerox Development Environment—expedites the integration of voice into workstation-based applications. Workstation program libraries implement the client programmer interface to the voice system. The voice control server associates each workstation with the physically adjacent Etherphone and interprets control requests accordingly. The physical distinction between Etherphones and workstations allows a variety of workstations to be accommodated without requiring additional voice hardware development.

The server software and the initial workstation software were developed in the Cedar programming environment [24]. More information on the equipment and protocols used in the Etherphone system, as well as the applications built to date, can be found in related papers [22, 23].

2.2 Some Applications of Recorded Voice

The desire to annotate documents with voice passages spurred the development of facilities for recorded voice in the Etherphone system. More specifically, users want to attach voice to any point in a document, to play that voice on request, to move or copy annotations readily and quickly, and to store such annotated documents in their host file systems and common file servers. Users want to share annotated documents with their colleagues, who are perhaps using different types of workstations with different file systems and document editors. Users also want the ability to edit voice annotations.

Figure 2 depicts two Cedar *viewers* (windows) involved in a voice-editing session. In the top viewer, the Tioga multimedia editor displays a document that includes two voice annotations, normally shown as cartoon balloons. One of the

ACM Transactions on Computer Systems, Vol. 6, No. 1, February 1988.

Managing Stored Voice in the Etherphone System

7

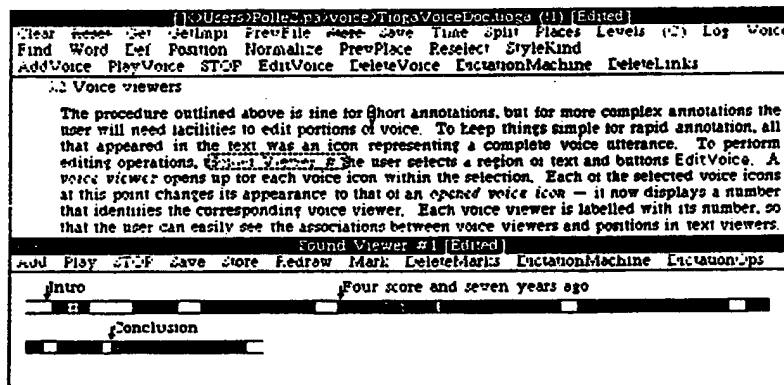


Fig. 2. An example of voice annotation and editing.

annotations has been "opened" to produce the lower viewer, which uses the alternating dark and light bars to represent intervals of voice and silence proportional to the lengths of the bars. The user employs conventional Tioga editing operations to delete, copy, or reorder sections of the original voice annotation. New spoken material can be recorded into selected locations by invoking special editor extensions. At any time, the user can listen to selected passages. Once edited, the result is then stored back as an amended annotation. This annotation remains with the document and with any copies that are made. The voice editor encourages the replacement and reorganization of recorded voice at the level of entire sentences or phrases. Finer-grained word or phoneme-level operations are discouraged for a variety of practical reasons explored in more depth in a related publication [1].

The Etherphone system also includes a voice mail prototype, which is another application of the voice annotation and editing facilities just described. The main difference between mail messages and ordinary annotated documents is that the messages are sent by a general mail-transport mechanism to lists of recipients, not manually filed by name in a filing system. Adding voice messaging to an existing environment is complex since mail is often read by a variety of mail programs running on a variety of systems. Most of these programs, both within our own and other interconnected networks, are not yet prepared to play or otherwise deal with recorded voice annotations.

Other applications that need to be able to record, modify, and store voice may employ substantially different user interfaces and storage models than those required of annotated documents. For example, we have found it useful to retain a number of prerecorded prompts, announcements, and attention getters in a simple directory. Another application that has been proposed would improve the ability to gather data during user studies, by recording the extemporaneous statements of users and synchronizing these recordings with time-stamped logs of other user actions. Additional applications, whose purposes and user-interface

8 • D. B. Terry and D. C. Swinehart

requirements we cannot know in advance, are expected to be developed in the future.

In general, adding such voice facilities to a diverse and complex software base presents challenging problems to the systems builder since much of the existing workstation and server software cannot be changed or extended. The voice-management facilities described in this paper were designed to make such applications less tedious to build, more robust, and easier to comprehend.

3. OPERATIONAL OVERVIEW

3.1 Voice Ropes

The implementation of facilities for recorded voice is somewhat involved, but the actions to be performed are conceptually quite simple. In looking for an application-independent abstraction to present to application programmers, it occurred to us that many of these actions closely resembled operations normally associated with text-string manipulation. The Cedar system provides a powerful text-string abstraction called a *rope* [24]. Cedar ropes are text strings of arbitrary length, represented as memory pointers, or references, to storage that is managed automatically by the Cedar system. When no references to a rope remain active, its storage is reclaimed. Cedar ropes are immutable: Once created, a rope's value will not change, so that ropes can be shared simply by sharing references. To create modified values, editing operations are supplied that create new ropes.

By analogy, in the Etherphone system, we refer to sequences of stored voice samples as *voice ropes*. A unique identifier, called a VoiceRopeID, is used instead of a memory pointer to identify each voice rope because, unlike Cedar ropes, voice ropes are persistent objects with a potentially long lifetime. To aid in sharing and to facilitate the use of voice by heterogeneous workstations, the storage for voice ropes, as well as the operations on them, is provided by a network service, the voice manager.

Clients refer to voice ropes solely by references, that is, by their unique VoiceRopeIDs. The voice manager places no restrictions on a client's use of voice ropes. Most uses involve embedding speech in some type of document, such as an annotated manuscript, program documentation, or electronic mail. The use of such embedded references to refer to voice, video, and other diverse types of information has been termed a *hypermedia system* [26].

From a client's perspective, a voice-annotated document should behave as though the voice were stored directly in the document's file rather than being included by reference. For example, once a voice message is sent using electronic mail, the author or another user should not be able to change the message's contents. For this reason, voice ropes, like Cedar ropes, are immutable. The recording and editing operations create new voice ropes; they do not modify existing ones.

3.2 Recording and Playing

To record or play a voice rope, a conversation is set up between the voice manager and an Etherphone. The main operations supported by the voice manager

ACM Transactions on Computer Systems, Vol. 6, No. 1, February 1988.

are as follows:

RECORD[conversationID] → voiceRopeID, requestID

Voice received by the server over the communication path defined by the given conversationID is stored and assigned a unique voiceRopeID; recording continues until a subsequent STOP operation is issued. The requestID identifies this operation in subsequent reports (see below).

PLAY[conversationID, voiceRopeID, interval] → requestID

The specified interval of the voice rope is transmitted over the given conversation. An interval denotes either the entire voice rope or a time-indexed portion of it at a resolution of about 1 ms.

STOP[conversationID]

Any recording or playing operations that are in progress or queued for the given conversation are immediately halted.

The RECORD and PLAY operations are performed asynchronously. That is, the RPC returns after the operation has been queued by the server. Queued operations are performed in order.

The voice manager generates event reports upon the start and completion of a queued operation. The requestID returned by each invocation is used to associate reports with specific operations. In particular, the voice manager makes the following call to all participants in a conversation to inform them of the status of various requested operations concerning that conversation:

REPORT[requestID, {started | finished | flushed}]

The requested operation has been started, successfully completed, or halted by a STOP operation.

Having reports flow from server to clients is conceptually similar to Clark's upcalls and accomplished in a similar manner [7].

3.3 Editing Support

Once recorded, voice ropes can be used in editing operations to produce new, immutable voice ropes. Several of the operations on Cedar ropes, such as producing substrings or concatenating existing strings, are directly applicable to voice. Their transliteration for voice ropes yields these functions:

CONCATENATE[voiceRopeID₁, voiceRopeID₂, ...] → voiceRopeID

Produces a new voice rope that is the concatenation of the given voice ropes.

SUBSTRING[voiceRopeID, interval] → voiceRopeID

Produces a new voice rope consisting of the specified interval of voiceRopeID₁.

REPLACE[voiceRopeID, interval, voiceRopeID₂] → voiceRopeID

Produces a new voice rope that is obtained by replacing the particular interval of voiceRopeID₁ with voiceRopeID₂. This is a composition of the CONCATENATE and SUBSTRING operations provided for efficiency and convenience.

LENGTH[voiceRopeID] → length

Returns the length of the given voice rope in milliseconds.

10 • D. B. Terry and D. C. Swinehart

One additional operation peculiar to voice ropes was provided to aid in editing:

DESCRIBE[voiceRopeID] → intervals

Returns a list of time intervals that denote the nonsilent *talkspurts* of the given voice rope. A *talkspurt* is defined to be any sequence of voice samples separated by some minimum amount of silence.

These operations, available via RPC calls to the voice manager, are intended for use by programmers. Applications that handle voice must employ these operations to construct the facilities visible to the end user.

3.4 Access Control

To ensure privacy, access-control lists govern who is permitted to play or edit particular voice ropes. Associated with each voice rope are two types of access: *play* access allows a client to play a voice passage, while *edit* access allows the client to use it in editing operations. Access-control lists may contain any number of individual or group names that are registered with the local name service, Grapevine [3]. The creator of a voice rope can change these access-control lists at any time by calling

PERMIT[voiceRopeID, players, editors]

Restricts access to the specified voice rope to the lists of players and editors.

Each newly created voice rope is given the default access controls specified by its creator. Typically, voice ropes are initially given unrestricted access or else restrict access to the voice rope's creator. Clients can later adjust permissions explicitly by calling **PERMIT**. For instance, a mail-sending program could routinely set the play-access-control list for a voice message to be the set of intended recipients.

3.5 Interests

The voice manager also provides operations for managing voice references. These operations provide a sort of directory for voice ropes. Although simple applications can use this directory as their means for naming and locating voice ropes, that is not its primary purpose. As with any storage system, unreferenced storage space should be reclaimed. With voice, or other voluminous media such as video, the need is particularly acute. Because voice ropes are shared by multiple users and multiple documents, manual management is impractical, and some form of garbage collection is required. In the Etherphone system, client code must assist with garbage collection by using the directory operations to express an *interest* in each referenced voice rope. The client operations are listed here, while a discussion of the rationale for this approach and a description of the underlying implementation is deferred to Section 4.5:

RETAIN[voiceRopeID, class, interest]

Registers an interest of the particular class in the given voice rope. The interest uniquely identifies a reference to the voice rope within the class. This operation is idempotent; successive calls with the same arguments register at most one interest in the given voice rope.

FORGET[voiceRopeID, class, interest]

Deregisters the specified interest and deletes the voice rope unless there are other interests for it.

LOOKUP[class, interest] → LIST OF voiceRopeID

Returns the unordered list of voice ropes associated with a particular interest.

Both interest and class are arbitrary text-string values. The form of the interest value is generally class specific. That is, each class controls its own name space of interests and may choose to use hierarchical, flat, or some other form of identifiers for its interest values. Clients are responsible for generating unique values for different interests within a class.

The class of an interest identifies the way in which voice ropes are being used by a particular application. For example, we use the class "FileAnnotation" to indicate that a document stored in a named file is annotated by a set of voice ropes; the interest field is the file name. The class "Message" indicates that an electronic-mail message incorporates recorded voice; in this case, the interest is the unique postmark supplied by the message system.

A combination of client workstation software and automatic collection methods must hide these interest operations from actual users. Client applications must always register interests in order to ensure retention of voice ropes to which they hold references. For some classes, clients must also explicitly FORGET their interests; for others, such as the "FileAnnotation" class, automatic methods described in Section 4.5 make this unnecessary.

4. DETAILED DESIGN DECISIONS


4.1 Voice-Manager Architecture

The voice manager uses a *voice file server* to store voice data. This server provides RECORD, PLAY, and STOP operations that are semantically similar to those described in Section 3.2, but operate on *voice files*. The more complex voice-rope editing and directory structures have been implemented as separate, higher level components, in part to make the voice facilities independent of the choice of the underlying file storage. Voice ropes in the current implementation are actually made up of pieces of one or more voice files, but could just as well be references to ordinary Cedar files, direct disk addresses, or address values from any other meaningful space.

The implementations of both voice rope editing and interest management depend on a simple but robust *database facility* that was developed for these purposes. Although the operations for editing voice ropes have been patterned after the Cedar rope package, a different underlying implementation was necessitated by the disparate characteristics of voice and text. Specifically, voice ropes are persistent, not transitory, values. Additionally, editing voice by actually copying the bytes, as is sometimes done for Cedar's ropes, is expensive since voice is voluminous. Thus, rather than rearranging the contents of voice files to edit them, the voice manager simply builds data structures using the facilities of the database package. Interests are registered in a similar database.

A *garbage collector* uses interests to reclaim voice storage that is no longer needed. Devising techniques for automatically collecting garbage in a distributed.

Fig. 3. Voice storage components.



heterogeneous environment was one of the most difficult problems faced in the design of the voice manager.

The components of the voice manager are logically layered as in Figure 3. Each is discussed in more detail in the following sections.

4.2 Voice File Server

A voice file server differs from conventional file servers [21] in that it must support the real-time requirements of voice. In particular, it must be able to maintain a sustained transfer rate of 64 kbits/s and should be able to support several such transfers simultaneously. There is no inherent reason why a general-purpose file server could not be extended to support these stringent real-time requirements. However, the file systems we had available at the onset of this project, having been optimized for different styles of access, could not, in fact, support them. At present, our file server is a special-purpose extension of Cedar's standard file system [24].

The Etherphone voice file server implements the operations needed to allocate, record, and play voice files that are named by unique identifiers, *VoiceFileIDs*. As previously mentioned, Etherphones encrypt voice as it is transmitted, using the DES encryption key associated with the current conversation. The voice file server simply stores the voice in its encrypted form. The voice rope implementation assumes responsibility for managing the encryption keys associated with various voice files. The stored voice is never decrypted except by an Etherphone when being played. Tables locating the boundaries between sound and silence are stored along with the voice to permit efficient execution of the DESCRIBE operation.

We will not describe the workings of the voice file server in greater detail. For the purposes of this paper, we assume that the reliable recording of voice files and the reliable playing of arbitrary queued sequences of voice file intervals are solved problems. We also presuppose the existence of relatively high-bandwidth network connections between the voice file server and Etherphones, a condition that is easy to meet in our local network environment.

4.3 Database Facilities

Voice ropes and interests rely on a simple database management system for their implementation. The requirements placed on the database system are not particularly stringent. It need only store immutable objects, provide basic query and update mechanisms, and support sharing among many client programs. There is no need for multiobject atomic updates, join operations, or comprehensive query languages. Any database system satisfying these modest requirements would suffice. Since such a system was not available in the Cedar environment, we developed a simple, robust database representation that is particularly well suited to voice ropes.

The database system stores each entry, a sequence of attributes expressed as key/value pairs, in a write-ahead log [10]. Unlike most database systems in which the data are logged only until they can be committed and written to a more permanent location, the log is itself the permanent source of data. Once logged, the data are never moved. To allow rapid queries or enumeration of database entries, B-tree indices [2] are built to map the values of one or more keys to the corresponding locations in the log file.

This log-based database package was inspired by similar methods found in the Walnut electronic-mail system [8] and recommended by Lampson [11]. In addition to managing the storage of voice ropes and interests, it has been used successfully for other data-management needs in the Etherphone system.

4.4 Voice-Rope Structure

The data structure representing a voice rope consists of a list of [VoiceFileID, key, interval] tuples. Additionally, an entry in the voice rope database contains attributes for the identifier, creator, access-control lists, and overall length of the voice rope. Thus, a typical database log entry for a voice rope is as follows:

```
VoiceRopeID: Terry.pa#575996078
Creator: Terry.pa
Length: 8000
PlayAccess: VoiceProject.t.pa
EditAccess: none
VoiceFileID: 235
Key: 17401121062B 10300460457B
Interval: 0 80000
```

A database index allows voice rope structures to be retrieved efficiently by VoiceRopeID; an index is also maintained on VoiceFileIDs, which is useful in garbage collection.

The database entry given above is for a simple voice rope consisting of a single interval within a single voice file. More complex voice ropes can be constructed using the editing operations presented in Section 3.3. For example, suppose two simple voice ropes, VR_1 and VR_2 , exist with the following structures:

```
 $VR_1 = \langle \text{VoiceFileID: } VF_1, \text{ Key: } K_1, \text{ Interval: [start: 0, length: 4000]} \rangle$ 
 $VR_2 = \langle \text{VoiceFileID: } VF_2, \text{ Key: } K_2, \text{ Interval: [start: 500, length: 2000]} \rangle.$ 
```

14 • D. B. Terry and D. C. Swinehart




Fig. 4. Structure of VR_3 after a REPLACE operation.

Then the operation

$\text{REPLACE}[\text{base: } VR_1, \text{interval: } [\text{start: } 1000, \text{length: } 1000], \text{with: } VR_2]$

produces a new voice rope, VR_3 , with the structure

$VR_3 = \langle \text{VoiceFileID: } VF_1, \text{Key: } K_1, \text{Interval: } [\text{start: } 0, \text{length: } 1000],$
 $\text{VoiceFileID: } VF_2, \text{Key: } K_2, \text{Interval: } [\text{start: } 500, \text{length: } 2000],$
 $\text{VoiceFileID: } VF_1, \text{Key: } K_1, \text{Interval: } [\text{start: } 2000, \text{length: } 2000] \rangle,$

as depicted in Figure 4.

To record a new voice rope, the voice manager calls on the voice file server to create a new voice file and store the voice arriving over a specified conversation as its contents. Once recording completes, a simple voice rope is added to the voice rope database to represent the complete voice file just recorded. The encryption key used to encode the conversation, and hence the voice file, is stored in the voice rope's entry. This key is carried along with the VoiceFileID during subsequent editing operations involving intervals of the voice rope. Independently enciphering small blocks of voice using ECB encryption [16] ensures that voice can be edited on millisecond-resolution boundaries while remaining encrypted. Note that the actual voice is neither moved nor copied once recorded in a voice file, even during editing.

When playing a voice rope, the voice manager retrieves the voice rope's structure from the database, checks access permissions against the caller's authenticated identity, distributes the encryption keys of the various intervals to the parties participating in the conversation, and calls upon the voice file server to play the intervals of the voice rope in the appropriate order. The voice file server provides sufficient buffering to support playing a queue of two or more voice file intervals without introducing gaps between the intervals.

The structure of voice ropes is kept "flat" to enhance playing performance. By having each voice rope refer directly to voice files, only a single database access is required to determine the voice rope's complete structure. An alternative design, more closely modeled on the Cedar rope abstraction, would store complex voice ropes as intervals of other voice ropes. In such a design, a voice rope would be expressed as a tree of other voice ropes with intervals of voice files at the leaves of the tree. This alternative design would reduce the work associated with

each editing operation, but increase the number of database accesses required to play a voice rope. The flat design was chosen because it improves playing behavior, and in practice, playing is much more frequent than editing. Moreover, it yields simpler and more compact data structures when used to represent small numbers of coarse-grained edits to voice. The alternative tree design would be a sensible approach within an environment optimized for a different distribution of usage patterns.

4.5 Storage Reclamation

4.5.1 Rationale for Interests. Once all voice ropes that reference a given voice file have been deleted, no voice rope will ever again refer to that voice file, so the voice file can be deleted as well. This condition is easily determined by a database query. The more difficult problem is deciding when voice ropes themselves can be reclaimed. The interest operations of Section 3.5 were included primarily to permit automatic reclamation of storage for voice ropes and their associated voice files.

From the client standpoint, the most straightforward method for garbage-collecting voice ropes would be periodically to examine all of the clients' storage for references to voice ropes, and then to collect any unreferenced ropes in a conventional sweep pass. Unfortunately, this sort of distributed garbage collector would be impossible to implement in our open, heterogeneous environment. We do not wish to restrict the uses clients make of voice ropes, or how and where clients store VoiceRopeIDs.

A common alternative is to provide a reference-counting scheme in which counters are used to determine the number of clients interested in a particular object. When an object's counter goes to zero, the object's storage can be reclaimed. The burden is placed on clients to increment and decrement the counts for the objects they are using.

The use of standard reference counting presents formidable problems in a distributed environment. Reference counts cannot be managed reliably unless an atomic transaction can be made to span all the activities involved in the creation and deletion of a reference. We consider this impractical to arrange, given our desire to accommodate a heterogeneous collection of participating systems. Without transactions, if the server or client fails in the process of incrementing or decrementing a reference count, the client may be left in an uncertain state regarding the outcome of the operation. In particular, client failure-recovery procedures might incorrectly repeat a reference-count operation. Furthermore, it is not always possible to arrange for a reference count to be decremented, such as when a reference-containing file residing on a noncollaborating server is deleted. Finally, reference counts are anonymous, giving no help in locating erroneous references.

Interests were designed to remedy the shortcomings associated with simple reference counts. In a way, interests represent a return to the full-scan method first proposed: Since an examination of the entire environment is not possible, clients are required to record their voice rope interests in a known place. The interests serve as proxies for the actual references for reclamation purposes.

Interests address two problems: how to retain voice ropes for which references exist, and how to determine when voice ropes can be reclaimed.

4.5.2 Retention of Voice Ropes. The use of interests to retain voice ropes is straightforward. Calling RETAIN adds an entry to the system's *interest database*, provided the entry is not already there, recording the supplied VoiceRopeID, class, and interest values along with the user's identification. The interest database includes indices permitting queries based on any of these attributes. Since a given entry appears in the database at most once, the RETAIN operation is idempotent. It can be safely retried in case of failures or uncertainty. The information stored in the interest database is sufficient to allow either human administrators or client applications to determine whether or not an interest is still valid.

4.5.3 Interest Invalidation. Determining when to invalidate interest entries is more involved. Some client programs can determine both when to RETAIN an interest and when to FORGET it. One example is the electronic-mail system that implements the "Message" class for which the interest design was first developed. When the mail system deletes an instance of a text message that contains voice references, it issues a FORGET for all of the associated voice ropes. Another example is the "SysNoises" class, a set of useful recorded announcements and sounds that system administrators manage manually. The FORGET implementation simply deletes the associated interest entry or entries from the interest database, ignoring requests to delete nonexistent entries to ensure that FORGET is also an idempotent operation. When every interest for a voice rope has been forgotten, the voice rope is vulnerable for deletion.

As we discovered when we began including voice ropes as annotations embedded in ordinary structured text documents, arranging for clients to issue FORGET actions at the necessary times is not always easy. Consider the following scenario: A user records a voice rope and embeds a reference to it in a document; an interest in the voice rope is then registered for the document. The user then copies this document from the workstation to a public file server and announces its existence in a message to interested parties. Several months later, the user deletes the file, without remembering it had voice annotations. Unless further actions are taken, the interest, and hence the referenced voice rope, will never be reclaimed.

Expecting an arbitrary file server to delete interests is not really reasonable. In the above scenario, one could argue that the file server should have issued the necessary FORGET operation when the file was deleted. This implies, however, that the software running on the file server could or should be modified to recognize the existence of files containing voice references, to understand their internal structures, and to take appropriate action. Many of the file servers in a typical network environment, including ours, cannot be so modified, either because they are old and written in some obscure programming language, or because they were purchased from outside vendors from whom source code is not available.

In this case, although interests cannot be explicitly deleted by the agents whose actions invalidate them, the voice manager can determine automatically when a

particular interest is no longer valid. Suppose a knowledgeable workstation client program issues the operation

RETAIN[vrID: voiceRopeID, class: "FileAnnotation", interest: "annotatedFile"]

for each voice rope referred to in a file as the file is moved from temporary workstation storage to the file named "annotatedFile" on a public file server. At any later time, standard directory operations can be used to determine whether that specific named instance of the file still exists; if not, the associated interest is no longer valid and can be deleted.

4.5.4 Garbage Collection. For automatically locating and removing outdated interests, the implementor of any interest class can register a procedure of the following type with the voice manager:

GARBAGE[voiceRopeID, interest] → {Yes | No}

Determines in a class-specific way whether or not the given interest still applies to the particular voice rope.

As an example, for the class "FileAnnotation," this procedure returns Yes if and only if the file instance identified by the interest parameter still exists on some file server. A "Time-out" class records an expiration time as its interest value; its GARBAGE procedure returns Yes when the time-out has expired.

An *interest verifier* periodically enumerates the database of interests and calls the class-specific GARBAGE procedure for each interest. If the procedure returns Yes, then the verifier calls FORGET[voiceRopeID, class, interest] to delete the interest from the database.

A garbage collector for voice ropes also runs periodically. For each voice rope in the database, the collector (1) deletes the voice rope if no interests exist that reference it, and (2) deletes voice files used by the voice rope if they are no longer a part of any other voice rope. This process refuses to collect voice ropes that are too young in order to prevent a newly created voice rope from being collected before a client has the opportunity to express an interest in it. This method will find all unreferenced voice ropes, including those for which no client ever expressed an interest and those that might have been orphaned due to system errors; otherwise, the actions of FORGET alone would be sufficient to eliminate unreferenced entries from the voice rope database. Note that, unlike a mark-and-sweep style garbage collector, these algorithms can be safely executed while the system is running and need not complete a full pass through the database in order to perform useful work.

The decision to protect voice ropes for a time after their creation was a pragmatic one. An earlier design required the client programmer to specify an interest, possibly of class "Time-out," for every voice rope created in order to force an explicit statement of the minimum conditions for deleting it. But, since many voice ropes are often created in the process of editing an annotation or voice message, this approach would generate unnecessary interests protecting these intermediate values. Instead, by choosing a minimum lifetime comparable to the interval between invocations of the garbage-collection procedures, we leave ample time for a properly functioning application to express an interest in the final result.

18 • D. B. Terry and D. C. Swinehart

In summary, garbage collection takes place on three levels. The voice manager deletes voice files when they are no longer referenced by voice ropes. Voice ropes are deleted if they are old enough and no interests exist for them. Interests are either explicitly forgotten by client applications or automatically deleted based on a class-specific test for validity.

5. EXPERIENCE AND EVALUATION

5.1 Current Usage

Approximately 50 Etherphones are in daily use in the Computer Science Laboratory at Xerox PARC. Our current voice file server runs on a Dorado [12], a 2-3 MIPS workstation developed at Xerox PARC, with a 300 Mbyte local disk. Thus, it has the capacity to store over 7 hours of recorded voice; the actual storage capacity depends on the amount of suppressed silence. Most of our user-level applications to date have been created in the Cedar environment [24], although limited functions have been provided for Interlisp and for stand-alone Etherphones. We have had a voice mail system running since 1984 and a prototype voice editor available for demonstrations and experimental use since the spring of 1986.

The current voice rope database contains approximately 2000 voice ropes with a total of 7000 segments referencing about 800 distinct voice files. Thus, the average voice rope consists of 3.5 segments, and an average voice file is used by 2.5 voice ropes. Half of the voice ropes have only a single segment, but some have over 20. Half the voice files are used by exactly 1 voice rope.

Voice ropes have an average length of 15.5 sec. Ninety-five percent are less than 45 sec in duration. The average segment size among edited voice ropes is 3 sec, which is in accordance with our preference for large-grained edits.

Fifty-one Mbytes of storage is used by the voice file server to store the collection of voice files whose aggregate length is 54 Mbytes. Thus, only about 5 percent in storage is saved by suppressing silence. Summing up the total lengths of all voice ropes yields 180 Mbytes, which indicates that a factor of 3 is gained through using the database to represent edits.

A majority of the voice ropes in the current database were created to test the voice manager itself or its applications during their development, and are therefore somewhat artificial. We would expect voice ropes representing annotations in actual production use to be somewhat longer on average, to consist of longer average segments, and to contain fewer edits. Presently, the interest database has not grown large enough to provide interesting statistics.

5.2 Voice-Annotated Documents

Manipulating stored voice solely by textual references, besides allowing efficient sharing and resource management, has made it easy to integrate voice into documents. For example, we were able to build a local voice mail system without changing the mail-transport protocols or servers. Also, annotated documents can be stored on conventional file servers that are not aware that the documents logically contain voice. The annotation applications utilized facilities already available in the Tioga editor for associating additional named properties with arbitrary locations within a document.

ACM Transactions on Computer Systems, Vol. 6, No. 1, February 1988.

Significant performance benefits accrued by having documents refer to voice that is stored remotely. Although most requests to record or play voice ropes are initiated from a workstation, the voice data are never received by the workstation; instead, they are transmitted directly to the associated Etherphone. A typical text document containing several voice ropes as annotations might occupy 30,000 bytes of storage, whereas inclusion of voice, assuming a minute of total commentary, would swell its size to 500,000 bytes or greater. In fact, scanned images are included in Tioga documents by value, and the resulting sizes do present storage and performance problems even for high-performance workstations.

Sharing documents and electronic messages that have been annotated with voice references, however, requires high-bandwidth network connectivity among the participants. Providing the applications described here among remote sites connected by limited-bandwidth channels would require additional mechanisms, such as those developed as part of DARPA's experimental multimedia mail project [18], which are beyond the scope of this paper.

5.3 Editing Voice

We gained considerable experience with the voice manager by building the Cedar voice-editing system described in Section 2.2 and illustrated in Figure 2. The small set of editing operations provided by the voice manager, including the specialized DESCRIBE operation that identifies the sound and silence intervals, has proved to be a sufficient base on which to build a complex voice editor and a dictation machine. However, to reduce traffic to the voice manager, the Cedar voice editor maintains its own data structures to represent the edited voice temporarily. That is, the voice editor ended up replicating much of the functionality of the voice manager, something we were trying to avoid. Only when a user elects to save the edited voice passage does the voice manager get called to perform the necessary operations. Given this situation, a better approach may have been to let clients simply pass the voice manager a complete voice rope specification that it could store in its database.

Experience with the initial implementation of the voice manager revealed that editing a voice passage invariably produced a set of "temporary" voice ropes used only in the construction of the finished result. These objects were eventually collected by the garbage collector and did no permanent harm, but did create additional work for the voice manager. To alleviate the problem somewhat, the voice manager's interface was changed slightly so that an interval could be supplied for any voice rope parameter in any operation. This substantially reduced the voice editor's use of the SUBSTRING operation.

Event reporting is important in allowing the voice editor to coordinate its visual feedback with the activities of the voice file server. In particular, the voice editor moves a cursor along the screen as a voice rope is being played (the gray marker below "score" in the voice displayed in Figure 2). A report indicating the playing of a particular voice passage has started or finished is essential to synchronize the movement of the cursor with the transmission of voice data.

Although the voice file server writes files on disk so that 1-sec segments can be continuously transferred, clients are allowed to edit voice ropes on 1-ms boundaries. The file server could not possibly play a voice rope in real time if it had to perform a disk seek every millisecond. Fortunately, users of the voice

20 • D. B. Terry and D. C. Swinehart

editor are encouraged to insert, delete, and rearrange voice passages at the granularity of a sentence or phrase rather than trying to modify individual words or phonemes [1]. Thus, in practice, one rarely sees segments of a voice rope that are less than several seconds in length.

5.4 Interests

The notion of grouping interests into classes and providing class-specific garbage-collection algorithms is a useful and workable concept. However, we are still groping with the details of how best to use these mechanisms. We have found several interest classes to be useful in Cedar.

The Cedar mail system automatically registers and deregisters interests of class "Message" as voice messages are saved and deleted by recipients. In addition, a "Time-out" class has been used to retract an interest automatically after a certain amount of time. For instance, when sending a voice message, a time-out of a week or two can be set by the sender to give recipients a chance to receive the message and register their own more permanent interests if so desired. Of course, problems can arise if a recipient is on vacation for a period of time longer than the time-out. For this reason, voice files are archived before being deleted from the server.

For annotated documents in Cedar, the workstation software detects when a file is copied from the local disk to a public file server; it then calls RETAIN to register the appropriate "FileAnnotation" interests for the public file. Having workstation software automatically register interests as a file is copied to a file server works remarkably well. However, some important operations are not covered by this approach: renaming a file on a file server or copying files between two file servers. We see no way to detect such operations except by modifying file-server software.

We have defined the "FileAnnotation" interest class such that its interest represents a publicly stored file name including the version number. With this scheme, interests must be reregistered for each new version of the file, that is, whenever a file is written to a public server. Unfortunately, the times that people want to annotate documents are precisely those times when the document is being updated often, so many interests are registered repeatedly. We rely on the garbage collector to get rid of old interests. An alternative would be to register a file without a version number, but that causes minor problems if voice is deleted from the file but the file itself remains in existence.

If we could obtain the cooperation of the file servers that store our documents, we could do a significantly better and less cumbersome job of managing interests. A collaborating file server could examine files for voice-rope references and execute the necessary RETAIN and FORGET operations as the files were stored, copied, and deleted. This would reduce or eliminate the burden on application programs. We consider this a promising area for additional research.

5.5 Reliability

The voice file server, voice manager, and voice control server were implemented so they could run on separate physical processors. That is, they all communicate among themselves and with voice clients using RPC. In practice, we run all three on the same Dorado. There is little to be gained by running them separately.

Managing Stored Voice in the Etherphone System • 21

Table I. Performance of Operations on Simple and Complex Voice Ropes (in milliseconds)

Operation	Time: simple	Time: complex
RECORD ¹	264-390 (318)	—
PLAY ¹	213-821 (394)	450-617 (563)
CONCATENATE	219-367 (279)	878-1276 (1105)
SUBSTRING	203-736 (373)	239-608 (345)
REPLACE	398-909 (444)	683-1109 (937)
LENGTH	33-72 ²	30
DESCRIBE	163	1432
RETAIN	285-681 (539)	"
FORGET	290-763 (514)	531-1024 (715)
LOOKUP	33 ³	"

¹ Time from beginning of RPC call to "started" report received by client.

² For voice ropes whose length is not known in the database.

³ Identical performance for simple and complex voice ropes.

since the voice file server cannot record or play voice files if the control server is down. Similarly, the voice manager cannot record or play voice ropes if the voice file server is down. Moreover, for all practical purposes, voice cannot be edited if the voice file server is down because users invariably need to listen to the voice passages they are editing.

Thus, availability is not adversely affected by having the voice manager and file server colocated with the control server. If this server crashes or is otherwise unavailable, then no operations can be performed on stored voice. For the most part, this is simply an inconvenience to users in the same way that the unavailability of conventional file servers is an inconvenience. In Cedar, the file servers containing the important system files, fonts, and documentation are replicated to improve their availability. In our prototype, we have not found it necessary to pay the cost to provide a highly available voice file server.

The one exception to this concerns voice interests. Clients often wish to register or deregister interests in voice ropes independently of playing the referenced voice. For example, a "FileAnnotation" interest is registered when a voice-annotated document is copied from a personal workstation to a public file server. A user should not be prevented from performing such a copy simply because the voice manager is unavailable. We have also observed that the interests for voice messages fail to get properly registered or deregistered if a person saves or deletes a voice message while the voice server is down. This has led us to contemplate writing a program that enumerates a person's mail database and checks that all voice messages have properly registered interests. The better solution is to make the voice-interest database highly available, at least for updates. Rather than fully replicating the database, we have designed a mechanism whereby operations to RETAIN or FORGET a voice interest are logged locally by a user's workstation if the database server is unavailable; the operations in this log will be retried when the workstation detects the voice manager has returned to operation.

5.6 Performance

The time performance of both the voice file server and the voice rope facilities easily meet the requirements of intended applications. Table I gives the measured performance of the various operations that can be performed on voice ropes.

22 • D. B. Terry and D. C. Swinehart

Table II. Breakdown of Time (in milliseconds)
Spent in the PLAY Operation for a Simple
Voice Rope

Sub-operation	Time
Communication (RPC) overhead:	18
Database lookup and access control:	18
Distribution of encryption keys:	137-745
Schedule playback with voice file server:	4
Time to receipt of "started" report:	36
TOTAL	213-821

Measurements were obtained for both a simple voice rope that contains a single 5-sec voice segment and a complex voice rope consisting of 10 5-sec segments obtained from 10 different voice files with different encryption keys. For these experiments, the client is a Dorado on a 3-Mbit Ethernet, while the server is a Dorado on a 1.5-Mbit Ethernet; the client and server are separated by a single gateway. Times are given in milliseconds; for operations with significant variance in measured times across several experiments, a range of times is given with the average in parentheses.

The time given for a RECORD or PLAY operation is the period from the client initiation of the RPC call until the receipt by the client of a file server report that recording or playing has started. Clearly, the complete time for the operation depends on the length of the voice rope being recorded or played. This time can be easily computed given a voice transmission rate of 64 kbits/sec.

We were quite surprised by the large variance in the time observed for playing a voice rope. Further examination of the PLAY operation indicated where the time is spent, as depicted in Table II. The variability is due to the time required to reliably distribute one or more encryption keys.

As expected, the cost of the DESCRIBE operation increases almost linearly with the length of the voice rope. For most of the other operations, the time depends heavily on the cost of database updates rather than on the size or complexity of the voice rope involved. However, the complexity of a voice rope has an indirect effect on the cost of writing that voice rope to the database. The database system maintains a B-tree index that keeps track of all voice ropes containing part of a given voice file. Thus, when adding a voice rope to the database, an update to this index is required for each segment of the voice rope. This explains why operations such as CONCATENATE and REPLACE are significantly more expensive for complex voice ropes.

We do not have current measurements of the server load during RECORD or PLAY. Early experience with the voice file server, however, indicated it can handle about eight simultaneous connections [22].

During the development of the voice manager, we have exercised the garbage-collection routines and verified they properly identify the objects to delete, but we have not actually allowed them to delete anything. Because of this, there are quite a large number of voice ropes, voice files, and interests in existence. Even so, recent measurements indicate the interest verifier makes a complete pass through the interest database in 13 sec. The voice rope garbage collector runs in

ACM Transactions on Computer Systems, Vol. 6, No. 1, February 1988.

140 sec, or approximately 80 ms per voice rope. Thus, the entire garbage-collection suite for a database of this magnitude can be executed in just a few minutes. Running the garbage collectors once per night, we could support a considerably larger population of voice ropes and interests than currently exist.

This paper has been concerned with the editing and management of recorded voice, dealing with operations whose performances must be compatible with human response times: subsecond response at a peak rate of several operations per second is more than adequate. The measurements reported herein confirm that the voice manager meets these requirements. Moreover, since network communication time is not a significant factor in the timings reported, the voice-rope facilities would work in a substantially more extensive and heterogeneous network environment than we have tried to date.

6. RELATED WORK

Several companies provide speech message systems that can be accessed from standard telephones; one of the earliest examples of this type of system was IBM's experimental Speech Filing System, which was operational in 1975 [9]. Certainly, the Etherphone system's facilities can be accessed from telephones, but that was not the driving application. We were interested in allowing voice to be integrated easily into a user's existing means of digital communications, rather than forcing users to learn a completely new system. The Sydis Information Manager provides workstation control over the recording, editing, and playing of voice as in the Etherphone system, but requires special workstations called VoiceStations [17]. Ruiz also developed a prototype voice system that integrates voice and data into some simple workstation applications; however, he did not address the important issues of sharing stored voice [19].

Maxemchuk's speech storage system [14] provided many of the same facilities for recording, editing, and playing voice as our voice file server. (Actually, he provided much more control over the playing of voice than we do, such as the ability to vary playback speeds or adjust silence intervals.) The division of function between a main computer and storage computer is also quite similar to the separation between our voice manager and voice file server. However, Maxemchuk's system edits voice using divide and join operations that modify the control sectors of stored voice messages. Our technique of building data structures that reference voice files better supports sharing by making voice ropes immutable and simplifies the requirements placed on the voice file server. For instance, our techniques are very amenable to write-once storage technologies such as optical disks.

Version Storage in the Swallow system [20] has many similar characteristics to our voice manager; that is, it manages immutable objects of various sizes. Also, its "structured version images" used for large objects are similar to the data structures used by the voice file server to describe voice files. Unlike the voice manager, Swallow provides no editing mechanisms or garbage collection, just read and write operations. It does, however, maintain histories to link together objects that are derived from one another and supports atomic operations on multiple objects.

The Diamond Document Store [25], like the Etherphone system, manages documents that contain various media elements by reference; it also allows documents to be shared among users by reference. Because the Diamond system does not allow documents stored outside the system to reference internally stored objects, a simple reference-count mechanism suffices for deallocating objects that reside in the Document Store but are not referenced by any document or document folder. The Etherphone system, on the other hand, strives to provide voice services that can be used along with other existing services, such as the Grapevine mail system [3] and Alpine file servers [6].

The Cambridge File Server [15] was perhaps the first network-accessible storage system to require clients to take explicit action to prevent files from being automatically garbage collected. In particular, it deletes files that are not accessible from server-maintained, but client-updated, "indices." Thus, these indices play much the same role as the voice manager's interest database.

Several methods for distributed garbage collection have been documented in the literature. For example, Liskov and Ladin have a scheme in which all sites that store references to other objects run a garbage collector locally and send information about nonlocal references to a "reference server" [13]. In some sense, their use of a reference server is similar to our use of registered interests, but much more limited. One interesting contribution they make is how to build a highly available reference server; we could use these techniques to build an interest server.

7. CONCLUSIONS

The facilities for managing stored voice in the Etherphone system were designed with the intent of moving voice data as little as possible. Once recorded in the voice file server, voice is never copied until a workstation sends a play request; at this point the voice is transmitted directly to an Etherphone. In particular, although workstations initiate most of the operations in the Etherphone system, there is little reason for them to receive the actual voice data since they have no way of playing them.

Maintaining voice on a publicly accessible server—the voice manager—facilitates sharing among various clients. Clients can freely share references to voice ropes without incurring the overhead of transmitting the voice itself. Because voice ropes are immutable, even though they are incorporated into documents by reference, they exhibit copy semantics.

To support efficient editing, a two-level storage hierarchy is employed: Voice ropes refer to intervals of voice files. A given voice rope can consist of intervals from several voice files, and a given voice file can be used by several voice ropes. A database stores the many-to-many relationships that exist between voice ropes and files. Editing operations simply create new voice ropes from old ones and add them to the database.

The editing operations provided by the voice manager are similar to those in the Cedar rope package. This is intentional so that programmers can manipulate voice in the ways to which they are accustomed to dealing with text. The basic facilities to support editing reside on a server; workstations are responsible for providing a user interface that is integrated with their programming environment.

Several aspects of the voice manager were designed to accommodate the heterogeneous nature of our environment. Providing a single implementation of the voice-rope facilities on a server obviates the need for each different workstation programming environment to provide its own implementation. Moreover, the only requirements placed on a workstation to make use of the voice services are that it have an associated Etherphone and an RPC implementation. In particular, workstations need not have direct hardware support for encryption or voice I/O.

The voice manager reduces the work generally associated with building voice applications by providing a convenient set of application-independent abstractions for stored voice. It makes no assumptions about the way clients make use of its services. This particularly impacted the design of the voice garbage collector.

Automatic reclamation of the storage occupied by unneeded voice ropes is done using a modified type of reference counting. Clients register interests in particular voice ropes. These interests are grouped into classes and can be invalidated according to a class-specific algorithm. For the most part, users of voice applications are not aware of how or when interests are registered since the application software handles this transparently.

The Etherphone system uses secure RPC for all control functions and DES encryption for transmitted voice. These ensure privacy of voice communication, which is important even in a research environment, although the Ethernet is inherently vulnerable to interception of information. Storing the voice in its encrypted form protects the voice on the server and also means that the voice need not be reencrypted when played. All in all, the voice system actually provides better security than most conventional file servers.

The Etherphone system has provided an environment in which to explore the management of voluminous, shared data among distributed and heterogeneous workstation clients. One could argue that compression of digitally recorded utterances would eliminate the need for special treatment in our current voice applications. Even with compression, however, there is a performance penalty in manipulating such large objects. More importantly, high-resolution scanned images and real-time digital video recordings will continue to be voluminous. We believe the techniques presented in this paper are applicable to and beneficial for the management of these media as well as to voice.

ACKNOWLEDGMENTS

The design of voice ropes evolved for several years, and many people contributed valuable suggestions, including Polle Zellweger, Stephen Ades, Luis Felipe Cabrera, and Larry Stewart. John Ousterhout designed and implemented the voice file server. Lia Adams built an early version of the voice mail system. Michael Schroeder suggested the use of interests for garbage collection of voice messages. Stephen Ades's implementation of a voice editor allowed us to get some experience with voice ropes. Severo Ornstein, Larry Stewart, and Dan Swinehart started the Etherphone project in 1982 and designed the Etherphone equipment. We are grateful to the many Cedar implementors who provided a wonderful environment for the Etherphone system.

26 • D. B. Terry and D. C. Swinehart

Margaret Butler, Alan Demers, Bob Hagmann, Jack Kent, Guy Steele, Polle Zellweger, and many anonymous reviewers contributed insightful suggestions, comments, and corrections. Subhana Menis, Bridget Scamporrino, Polle Zellweger, and Rick Beach provided invaluable assistance with the composition and printing.

REFERENCES

1. ADES, S., AND SWINEHART, D. C. Voice annotation and editing in a workstation environment. In *Proceedings of AVIOS Voice Applications '86* (Alexandria, Va., Sept. 1986). Avios, 1986, pp. 13-28.
2. BAYER, R., AND MCCREIGHT, E. Organization and maintenance of large ordered indexes. *Acta Inf.* 1, 3 (1972), 173-189.
3. BIRRELL, A. D., LEVIN, R., NEEDHAM, R. M., AND SCHROEDER, M. D. Grapevine: An exercise in distributed computing. *Commun. ACM* 25, 4 (Apr. 1982), 260-274.
4. BIRRELL, A. D. Secure communication using remote procedure calls. *ACM Trans. Comput. Syst.* 3, 1 (Feb. 1985), 1-14.
5. BIRRELL, A. D., AND NELSON, B. J. Implementing remote procedure calls. *ACM Trans. Comput. Syst.* 2, 1 (Feb. 1984), 39-59.
6. BROWN, M. R., KOLLING, K. N., AND TAFT, E. A. The Alpine file system. *ACM Trans. Comput. Syst.* 3, 4 (Nov. 1985), 261-293.
7. CLARK, D. D. The structuring of systems using upcalls. In *Proceedings of the 10th ACM Symposium on Operating Systems Principles* (Orcas Island, Wash., Dec. 1985). ACM, New York, 1985, pp. 171-180.
8. DONAHUE, J., AND ORR, W.-S. Walnut: Storing electronic mail in a database. Tech. Rep. CSL-85-9, Xerox Palo Alto Research Center, Palo Alto, Calif., Nov. 1985.
9. GOULD, J. D., AND BOIES, S. J. Speech filing—An office system for principles. *IBM Syst. J.* 23, 1 (Jan. 1984), 65-81.
10. GRAY, J. N. Notes on database operating systems. In *Operating Systems: An Advanced Course*, R. Bayer et al., Eds. Springer-Verlag, New York, 1978, pp. 393-481.
11. LAMPSON, B. W. Hints for computer system design. In *Proceedings of the 9th Symposium on Operating Systems Principles* (Bretton Woods, N.H., Oct. 1983). ACM, New York, 1983, pp. 33-48.
12. LAMPSON, B. W., AND PIER, K. A. A processor for a high-performance personal computer. In *Proceedings of the 7th ACM Symposium on Computer Architecture* (La Baule, France, May 1980). ACM, New York, 1980, pp. 146-160.
13. LISKOV, B., AND LADIN, R. Highly-available distributed services and fault-tolerant distributed garbage collection. In *Proceedings of the ACM Symposium on Principles of Distributed Computing* (Calgary, Alberta, Canada, Aug. 1986). ACM, New York, 1986, pp. 29-39.
14. MAXEMCHUK, N. An experimental speech storage and editing facility. *Bell Syst. Tech. J.* 59, 8 (Oct. 1980), 1383-1395.
15. MITCHELL, J. G., AND DION, J. A comparison of two network-based file servers. *Commun. ACM* 25, 4 (Apr. 1982), 233-245.
16. NATIONAL BUREAU OF STANDARDS. Data encryption standard. Federal Information Processing Standard (FIPS) Publ. 46, U.S. Department of Commerce, Washington, D.C., Jan. 1977.
17. NICHOLSON, R. Integrating voice in the office world. *BYTE* 8, 12 (Dec. 1983), 177-184.
18. REYNOLDS, J. K., POSTEL, J. B., KATZ, A. R., FINN, G. G., AND DESCHON, A. L. The DARPA experimental multimedia mail system. *Computer* 18, 10 (Oct. 1985), 82-89.
19. RUIZ, A. Voice and telephony applications for the office workstation. In *Proceedings of the 1st International Conference on Computer Workstations* (San Jose, Calif., Nov.) 1985, pp. 159-163.
20. SVOBODOVA, L. A reliable object-oriented data repository for a distributed computer system. In *Proceedings of the 8th ACM Symposium on Operating Systems Principles* (Pacific Grove, Calif., Dec. 1981). ACM, New York, 1981, pp. 47-58.
21. SVOBODOVA, L. File servers for network-based distributed systems. *ACM Comput. Surv.* 16, 4 (Dec. 1984), 353-398.

ACM Transactions on Computer Systems, Vol. 6, No. 1, February 1988.

Managing Stored Voice in the Etherphone System • 27

22. SWINEHART, D. C., STEWART, L. C., AND ORNSTEIN, S. M. Adding voice to an office computer network. In *Proceedings of the IEEE GlobeCom '83* (San Diego, Calif., Nov. 1983). IEEE Press, New York, 1983. (Also Tech. Rep. CSL-83-8, Xerox Palo Alto Research Center, Palo Alto, Calif., Feb. 1984.)
23. SWINEHART, D. C., TERRY, D. B., AND ZELLWEGER, P. T. An experimental environment for voice system development. *IEEE Off. Knowl. Eng. Newsl.* 1, 1 (Feb. 1987), 39-48.
24. SWINEHART, D., ZELLWEGER, P., BEACH, R., AND HAGMANN, R. A structural view of the Cedar programming environment. *ACM Trans. Program. Lang. Syst.* 8, 4 (Oct. 1986), 419-490.
25. THOMAS, R. H., FORSDICK, H. C., CROWLEY, T. R., SCHAAF, R. W., TOMLINSIN, R. S., TRAVERS, V. M., AND ROBERTSON, G. G. Diamond: A multimedia message system built on a distributed architecture. *Computer* 18, 12 (Dec. 1985), 65-78.
26. YANKOLOVICH, N., MEYROWITZ, N., AND VAN DAM, A. Reading and writing the electronic book. *Computer* 18, 10 (Oct. 1985), 15-30.

Received May 1987; revised August 1987; accepted September 1987

THIS PAGE BLANK (USPTO)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☒ FADED TEXT OR DRAWING
- ☒ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (USPTO)